Computer Systems, Inc.
MERCURY
Challenges Drive Innovation

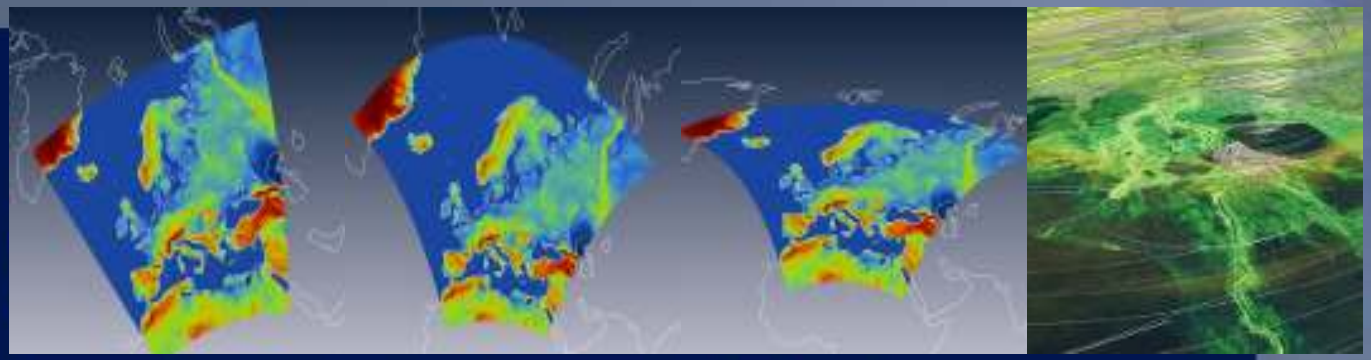

Flexible Visualization Environment for Post-Processing of Simulation Data from Numerical Earth System Models

Ultravis08
Austin,TX November 17th 2008
Patrick BARTHELEMY

Project Description

- **Customer**
  - DKRZ : Deutsches Klimarechenzentrum
  - www.dkrz.de
  - ***The German High Performance Computing Center for Climate and Earth System Research***
- **Project**
  - Visualization of time dependent data related to Climate & Earth observation.

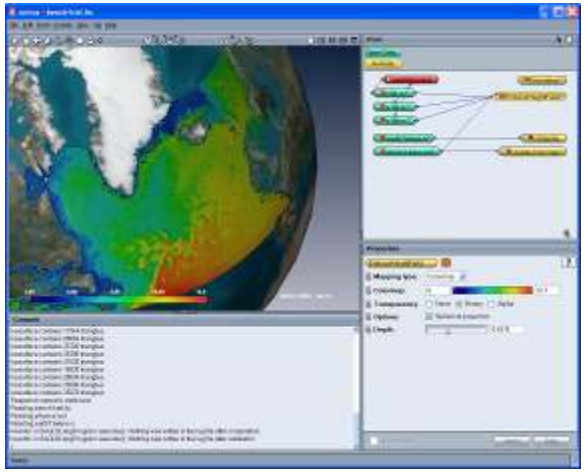

Requirements

- **Innovative solution to fulfill current and anticipated needs for climate data visualization**
  - Enable analysis of very large datasets, not restricted by available memory, even aggregated cluster memory
  - Provide effective support of climate data including time series datasets
  - Support a variable number of local or remote users
  - Enable desktop users to complement their local resources by exploiting the server capabilities
  - Support VR display systems and immersive interaction
  - Support cost-effective hardware, easily upgradeable

Solution

- **Avizo visualization framework**
  - multi-resolution data management and visualization streaming
  - parallel rendering with compositing
  - distributed data management
  - extensive use of GPU for data post-processing and representations
  - flexible remote visualization offering scalability to distant users
  - flexible display configurations and support for immersive VR with wide range of interaction devices
- **HP Scalable Visualization Array (HP SVA)**
  - Linux Red Hat Enterprise X_64 for cluster nodes
  - Windows XP or Linux Red Hat Enterprise X_64 for client
- **HP StorageWorks Scalable File Share high-performance storage system (HP SFS).**

DKRZ Visualization Server

- **Hardware**
  - 8 Vis Nodes
    - HP XW 9400 Workstation
    - 2 Dual Core Opteron64 2.6 GHz
    - 32 GB RAM
    - 2 NVidia FX 5500 Graphics with 1 GB RAM
    - InfiniBand HCA
  - 1 SMP Vis Node
    - *HP ProLiant DL585 Server*
    - 4 Dual Core Opteron64 2.4 GHz
    - 128 GB RAM
    - 2 NVidia FX 5500 Graphics with 1 GB RAM (via NVidia QuadroPlex)
    - InfiniBand HCA
  - 1 System Management Node
    - *HP ProLiant DL 385 Server*
    - 2 Dual Core Opteron 2.2 GHz
    - 4 GB RAM
    - InfiniBand HCA
  - HP SFS (Scalable File Share) 72 TByte brutto / 48 TByte netto
    - 2 Meta Data Server (MDS)
    - 6 Object Storage Targets (OST)
    - Infiniband HCAs
- **Network**
  - Gigabit Ethernet Network
  - System Management Network
  - 10 Gigabit Ethernet Uplink
  - InfiniBand Network with Voltaire 24 port switch (4x, single data rate: 10 GBit/s)

Data

- **Simulation results are stored using NetCDF CF 1.0 format**
- **Time dependent data**
  - Rectilinear, curvilinear 2D/3D grid
  - Scalar information for temperature, pressure…
  - Flow information for velocity, wind, current…
- **Size**
  - 3d datasets from single values can contain more then 100 Million data points.
  - Many of these values can be stored at any time step in a single file.
  - Several of these values must be visualized at the same time.
  - Very long time series of multi value datasets must be animated

Avizo Cluster Visualization

- **Rendering Farm**
  - Application distribute rendering on the cluster
  - Cluster compositing
    - Each GPU generates an image
    - All GPUs are used to make composition of images
  - Intra-cluster communication
    - Infiniband/MPI is used to communicate between nodes
- **Data Access**
  - Each node access data in parallel
    - Several I/O parallel threads are running on each node
  - Location of data
    - SFS
- **Remote Rendering**
  - Client Server communication
    - Ethernet & TCP/IP protocol
    - Low bandwidth requirements
  - Compression & Filtering
    - To reduce network traffic images are filtered & compressed
    - All GPUs are used to compress and filter
  - Client Decompression
    - Client GPU/CPU decompress the image coming from the cluster
  - Control of Image Quality
    - Client can control quality of compression to reduce bandwidth

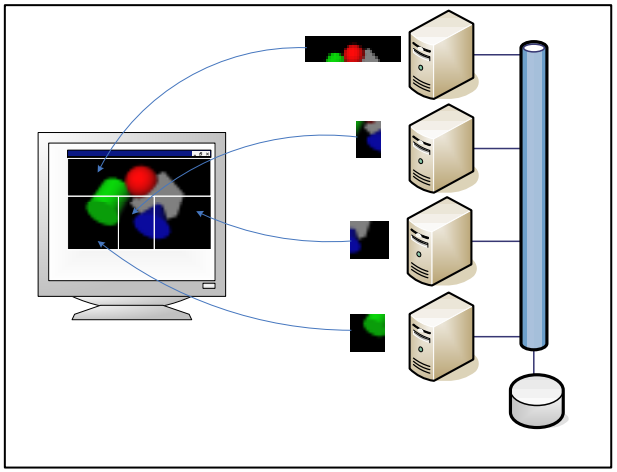

Tile Compositing

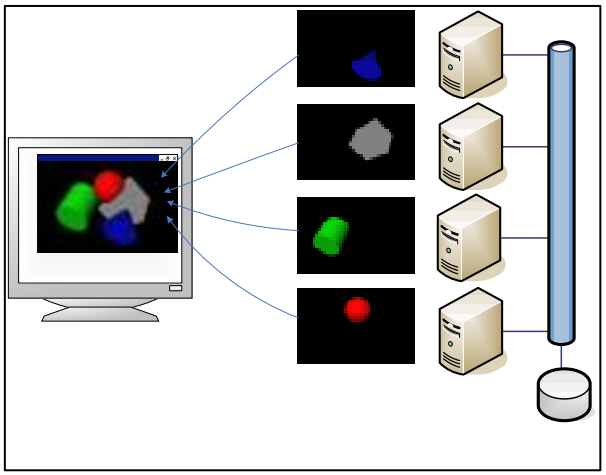

Depth Compositing

Avizo Cluster Visualization

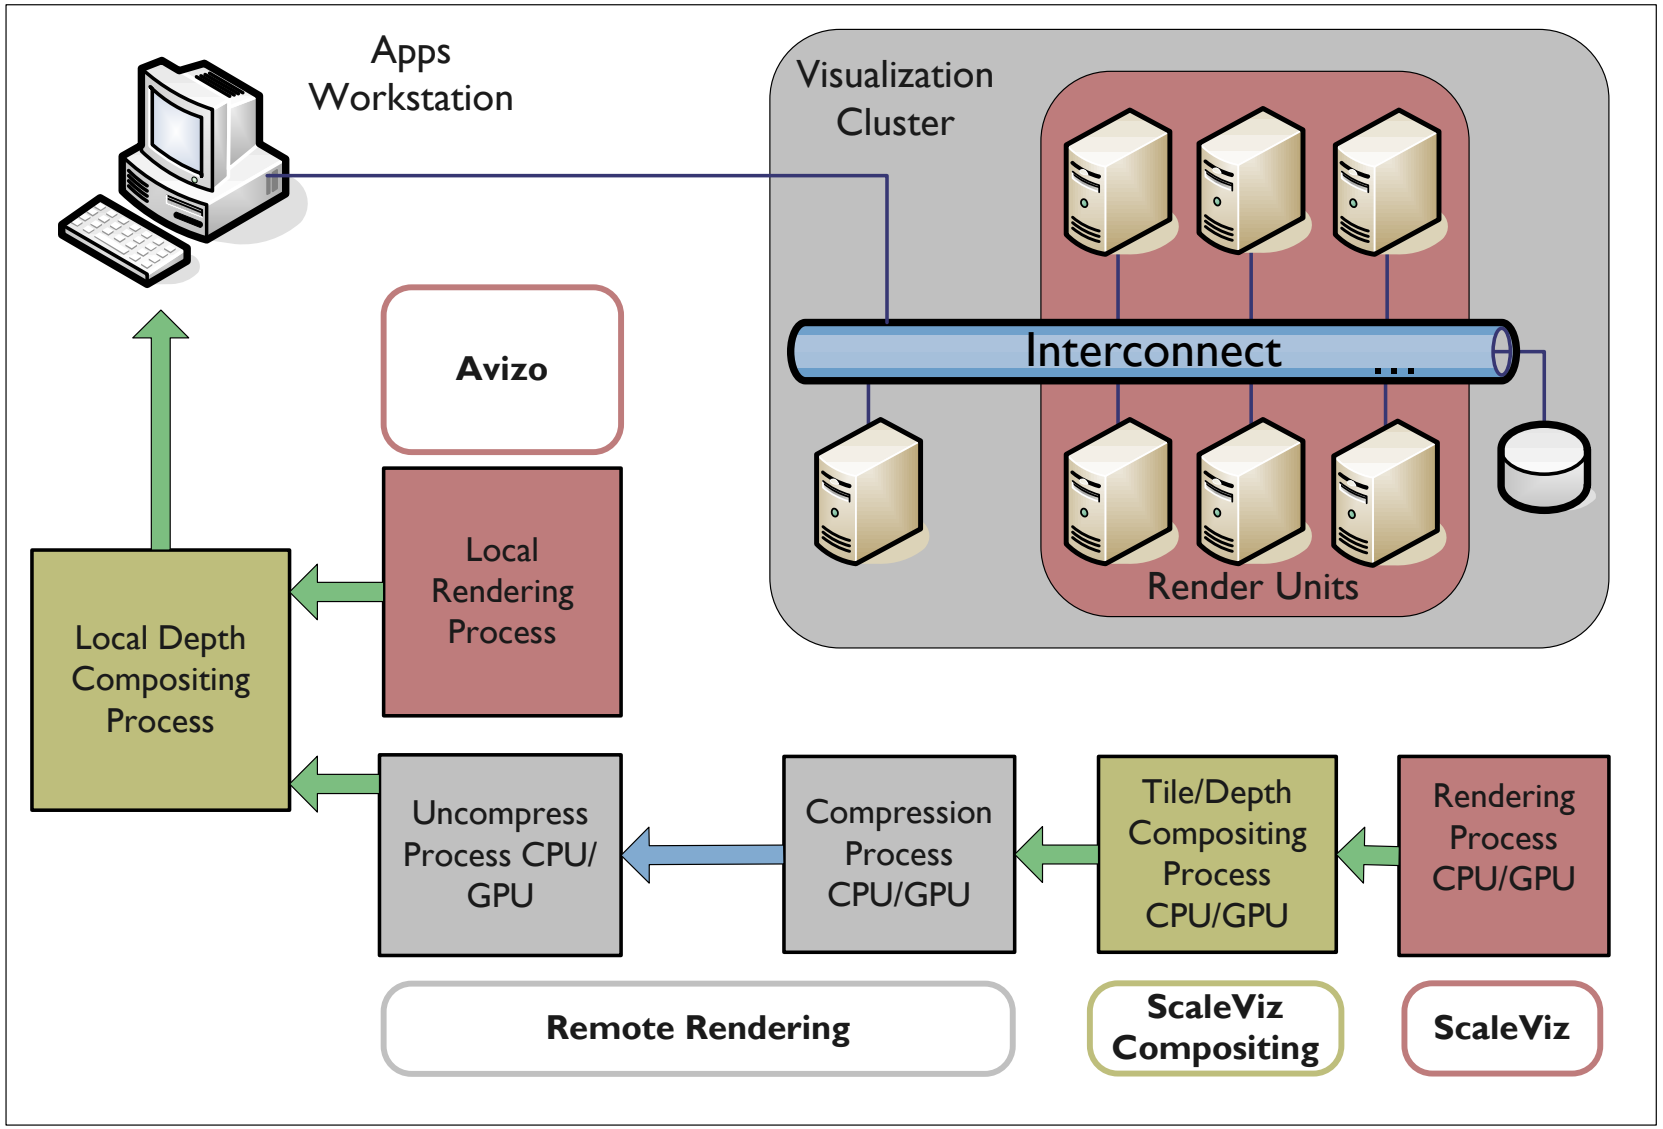

Avizo and Large Data Management

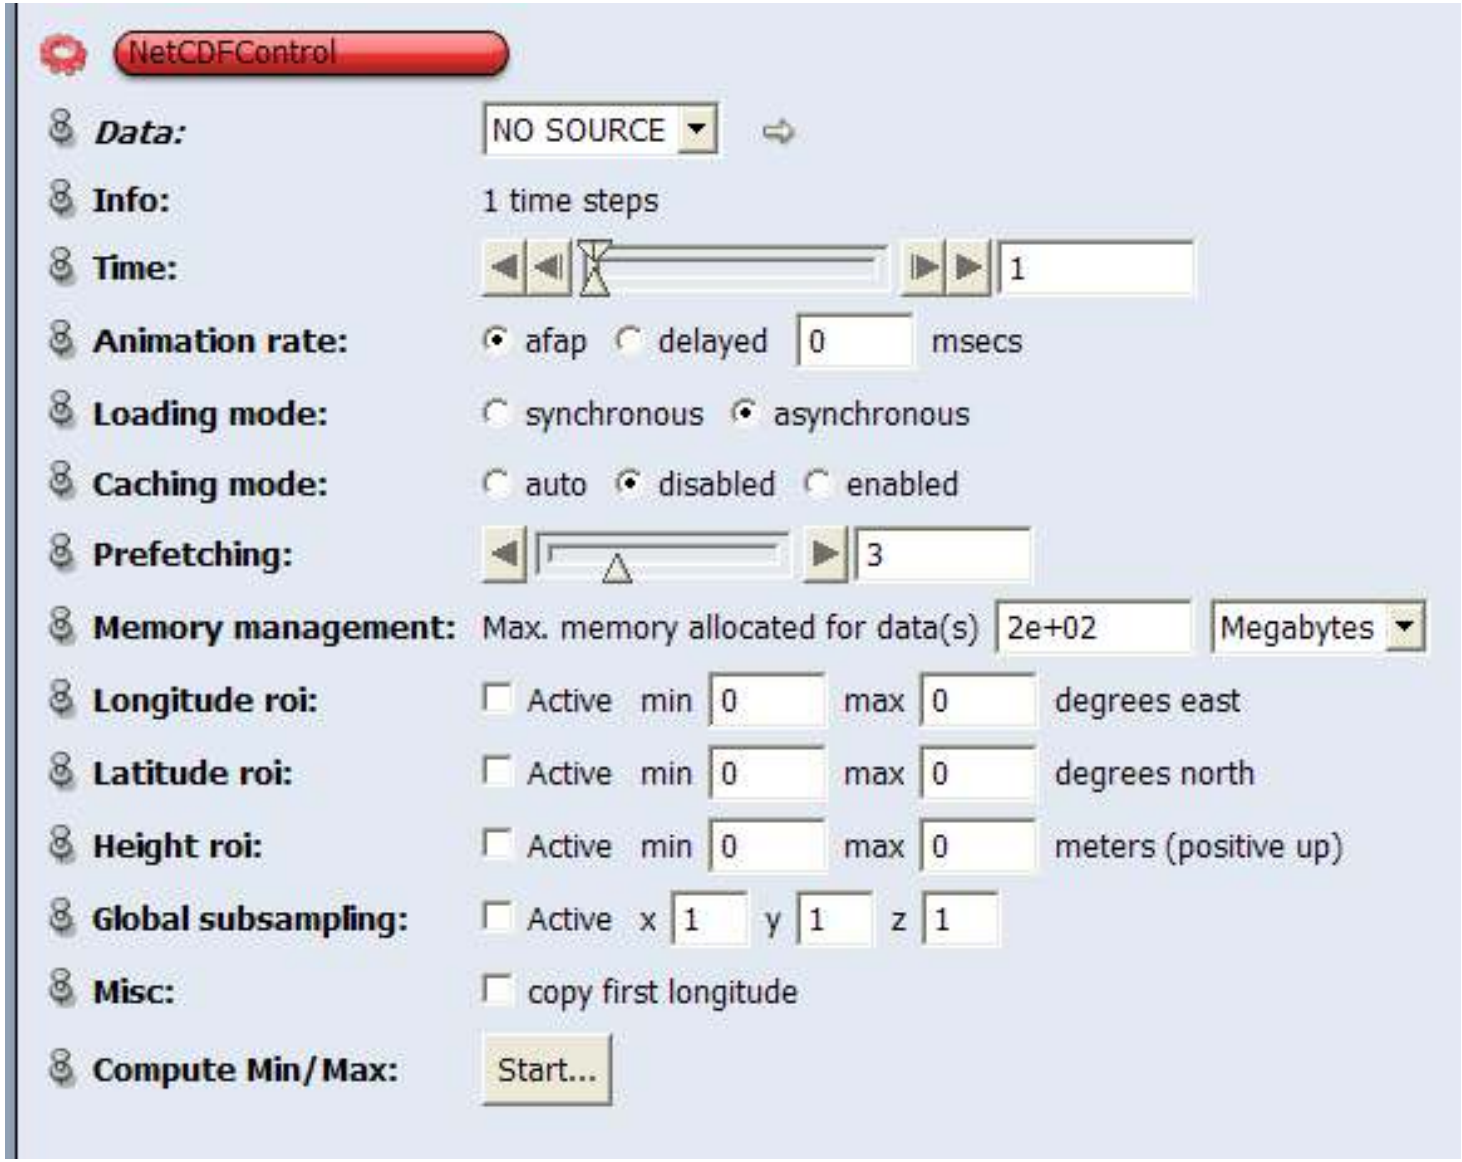

Solution in action

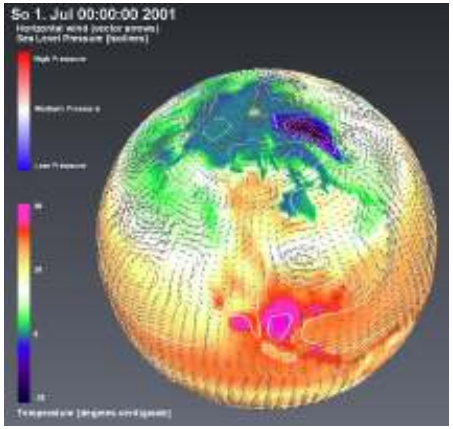

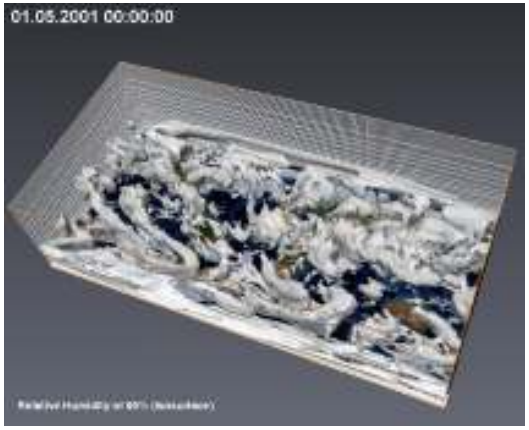

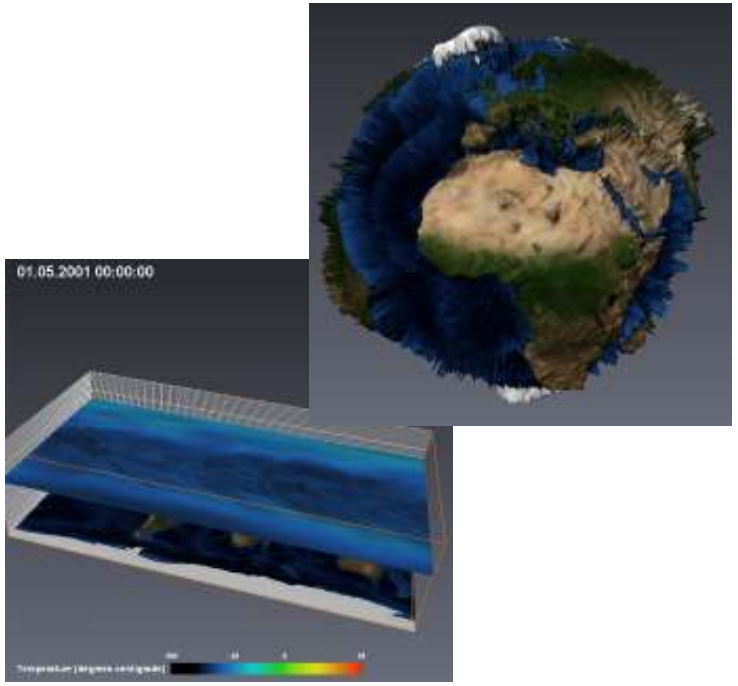

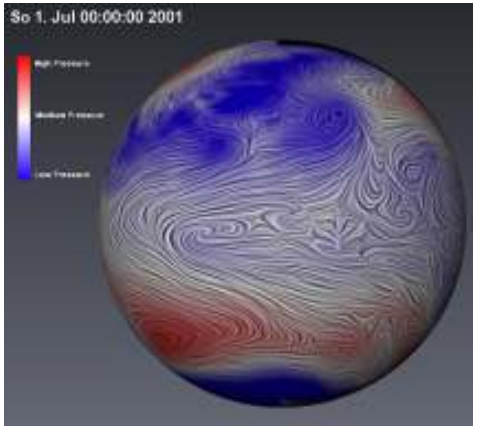

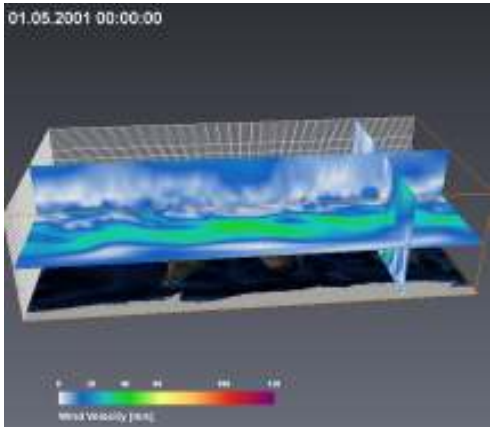

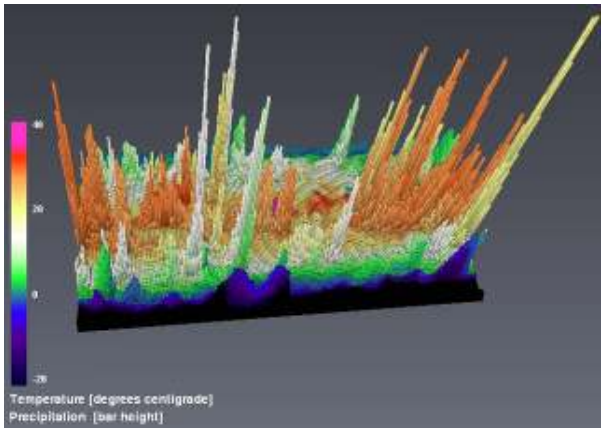

More information

- **Visualization Sciences Group**
  - http://3dviz.mc.com
    - http://www.avizo3d.com
    - http://3dviz.mc.com/products/avizo_XGreen_overview.asp

Thank you

3d_info@mc.com

pbarthel@mc.com